‘Civic Digital Visibility First’ - A Conditional Welcome to Our Emerging Civic AI, and a Call to Ban Democratically-Uncertified AI from Government Use

Christopher Hunt Robertson, M.Ed.

(Ohio Academy of History Speech - Marietta College - March 21, 2026)

The federal government and many state and local government offices are now actively encouraging the use of AI in government. That may be well-intentioned but it is risky.

Currently, about 85 percent of government offices rely on a single word processing software. Let’s imagine a hypothetical scenario where a company chose to introduce a new product: “*Civic Word: Now with AI that can prioritize government projects!*” As a manager, I think, “I can call my staff together for a long, heated debate about our projects, or I can just push a button and let Civic Word rank them for us.” At that moment, a democratic decision that once required human debate quietly becomes a software feature.

American democratic government - at the federal, state, and local levels - should reject all AI from government settings by default until that AI has been independently certified as compliant with democratic principles. We must not let ‘optimizing’ AI slip into our government sector on the strength of efficiency claims alone - Trojan Horses quietly treated as if they were Civic AI simply because they are deployed in civic settings.

We cannot jeopardize 2500 years of democratic evolution.

OUR RAPIDLY GROWING PROBLEM

Many assume that democracy is permanent - as if once established, it could only move forward, never backward. But we know that the historical record is not a straight line of democratic progress. It is a record of gains and losses - of republics that rose, stabilized for a while, and then often weakened or fell.

In recent years, we have watched signs of that weakening spread here in the United States and across the globe: eroding trust, captured institutions, elections held - but not believed.

These pressures did not begin with artificial intelligence. But AI - by neglect or by intent - can now accelerate these pressures at machine speed. It can deepen fractures faster than our civic efforts can repair them.

At the heart of this danger lies something very simple: democracies cannot govern what they cannot see.

So today I want to suggest that our central civic crisis is not a crisis of moral darkness. It is the fragmentation of the moral light that we still possess. But we're losing our ability to see clearly enough, quickly enough, and together enough to govern responsibly.

So how can democracies see clearly enough to govern responsibly in the digital age?

We must first understand the magnitude of the challenge. We can begin by looking at our own headlines. They reveal a pattern: how our individual and collective consciences are strained under a constant flood of crisis and confusion - and how that strain is beginning to erode the historically evolved civic processes on which democracy depends.

The pattern moves outward in six stages. Together these form what I call the *Six Rings of Digital Invisibility and Civic Destabilization*. Each ring widens the gap between democratic responsibility and democratic sight.

Ring One - Moral Overload

Modern citizens now absorb more suffering, outrage, and breaking news events in a single week than many earlier generations encountered in a year: war, climate disruption, political fracture, public heath crises, economic precarity, algorithmically amplified conflict.

The result is not moral indifference; it is moral saturation. Our sense of *ought* is triggered again and again - without resolution and without rest. People do not become immoral under that load; they become overwhelmed. Deliberation gives way to exhaustion. Outrage and withdrawal alternate. And when conscience is constantly triggered but rarely resolved, it slowly loses its capacity to guide.

Ring Two - Moral and Civic Invisibility

AI now operates inside financial markets at speeds regulators describe as “too fast to stop.” Financial deepfakes have already moved markets with false announcements. Real harm occurs - but it unfolds across platforms, jurisdictions, and milliseconds. By the time we see the consequences, the chain that produced them is no longer visible.

For most of human history, harm was nearby. Suffering had faces. Responsibility could be traced. Today, harm increasingly occurs at a distance - mediated by platforms and fragmented across systems. People may sincerely wish to do no harm and yet participate, unknowingly, in systems that cause it. Conscience does not disappear here - but it loses its bearings.

Ring Three - Institutional Blindness

Modern institutions increasingly rely on systems they cannot fully see into. Courts, hospitals, regulators, election systems, and public-benefits programs now depend on digital infrastructures whose inner workings remain opaque.

We read about AI-assisted cyberattacks generating constantly mutating malware. We see attacks on hospitals, airports, and water systems. Government agencies deploy algorithmic tools whose logic even their operators cannot fully explain. Authority remains but comprehension thins. Institutions exercise power with diminishing visibility into how outcomes are produced. And when systems act beyond meaningful human oversight, legitimacy begins to erode.

Ring Four - Amplification Without Judgment

Deepfake videos circulate before verification can catch up.
AI systems generate thousands of message variations to discover which narratives provoke the strongest emotional response.
A growing share of political content online is synthetic or heavily manipulated.

Our digital systems are extraordinarily good at scaling what travels fastest - outrage, fear, simplified stories. But they are not designed to distinguish wisdom from distortion. When synthetic media becomes indistinguishable from authentic recording, shared reality begins to fracture. Democratic deliberation cannot function if it cannot discern.

Ring Five - Governance Under Strain

Across the United States and the world, governments are scrambling to respond to AI's public impact. Hundreds of legislative proposals appear. Executives invoke competitiveness. Legislatures demand oversight. Courts attempt to referee technical disputes with only partial visibility.

Leaders must make consequential decisions without a stable shared picture of what is actually happening. Government shifts from stewardship to permanent crisis response. Authority reacts to pressure rather than acting from understanding.

Ring Six - Cascading System Consequences

A ransomware attack in one country shuts down hospitals in another. Synthetic narratives created by small groups are amplified by foreign actors, domestic factions, and automated systems. A single misperception can move from social media to financial markets, to politics, to street-level conflict.

The rings reinforce one another: overload - invisibility - institutional blindness - amplification without judgment - governance under strain. By the time the damage becomes visible, it is already widespread and politically difficult to repair. Consequences now arrive faster than responsibility can be traced.

The Diagnosis

The *Six Rings* are not a prophecy of democratic collapse. They are a diagnosis of what our own headlines already describe: a progressive breakdown in our shared ability to see, interpret, and coordinate responsibly in a world moving at digital speed.

The problem is not a sudden loss of intelligence or goodwill. It is something more structural. Our civic environment is now moving faster than human perception was built to manage. We face a growing mismatch between moral responsibility and perceptual capacity.

And if this crisis is a collapse of civic visibility, then the central question becomes this: What could help our democracy see again - together - without giving our moral authority to machines?

Perhaps part of the response lies in something unexpected. If AI-acceleration is amplifying our problems, perhaps the development of democratically-certified Civic AI could become part of our solution - as a new instrument of sight.

Franklin's Experiment

It may help to remember an early American experiment. During a storm, Benjamin Franklin flew a kite carrying a key into a thundercloud and revealed the invisible electrical forces moving through the sky. He made an unseen force visible to human understanding. That new visibility soon produced a practical invention: the lightning rod - a tool of protection.

Properly designed, Civic AI could serve a similar function for democracy today. Imagine Civic AI acting like Franklin's kite and key for the digital age - drawing down hidden, harmful digital currents that affect civic life and making them visible to human discernment, while never becoming an instrument of power itself.

Civic AI can make democracies faster at seeing. Every moment saved in identifying hazards, tracing patterns, or assembling information at machine speed becomes one more moment available for human deliberation. More time. A clearer record. Better judgment.

But when designing Civic AI, we must ensure that it does not damage our democracy more than the problems it is meant to address.

A PROPOSED SOLUTION: DEMOCRATICALLY-CERTIFIED CIVIC AI

Democracies are not designed to maximize efficiency. They are designed to limit the concentration of power. And in the digital age, that means limiting the concentration of agentic capacity - systems that can act, steer, and scale influence.

AI is now being rapidly integrated into federal government, as well as state and local governments. The problem I see is this: AI does not become Civic AI simply because it is deployed in civic contexts. I propose that AI be banned from government use until it is proven to be democratically-certified Civic AI.

But before defining Civic AI, we must consider several fundamental questions:

- Under what constitutional conditions may new power approach public authority at all?
- Will civic visibility remain shared public infrastructure?
- Will civic visibility become concentrated in systems and institutions that see more than citizens do?

We don't need Civic AI to tell us what is right. We need it to tell us what is real. And so, in the years ahead, I suspect one principle will matter most: democratic illumination.

Not capacity. Not agency. Democratic illumination.

Moral visibility is the structured illumination of civic conditions where human responsibility is at stake - without directing, ranking, or resolving that responsibility. It makes consequences visible: who is affected, how, and over what time frame. It reveals patterns and chains of responsibility that might otherwise disappear inside digital complexity, and it gives human conscience something honest to look at. It does not preach or assign virtue, rank citizens, or engineer agreement. Moral visibility does not tell us what to believe; it simply makes the stakes harder to ignore.

First sight, then deliberation. Democracies already understand this sequence. Courts, legislatures, and voters seek information and clarity before acting. And yet, in many current discussions of AI, that sequence is quietly reversed. We begin with persuasion - optimization - mobilization. And visibility somehow becomes an afterthought.

Democracy works best in the opposite order: First sight, then deliberation. Because democracies cannot govern what they cannot see.

A properly ordered civic process might look like this:

First, visibility - illumination of facts, hazards, and patterns.

Second, analysis on request - synthesis, comparison, explanation.

Third, human deliberation - citizens, journalists, courts, and legislatures weighing competing values.

Fourth, persuasion and activism - if chosen.

And finally, law and enforcement - if authorized.

Visibility does not compete with civic action. It makes civic action possible.

The premise of my *Civic Digital Visibility First* principle is simple. Treat Civic AI as visibility infrastructure - as digital lighthouses that warn us of digital hazards. Under this approach, a Civic AI system may illuminate hazards and inequities, reveal patterns and correlations, and provide analysis when requested.

But it must always stop short of the final sentence: "Therefore, you *should* do X." That final *should* belongs to human beings who must bear the moral and civic consequences of their choices.

Moral visibility is not a side effect here. It is the entire purpose. Civic AI should be a witness, not a sovereign. To maintain that boundary, we need explicit design constraints.

My work proposes what I call the *Fifty-Gates Democratic Permissions Threshold*. Each gate represents a democratic principle and acts as a binary (yes/no) constitutional test for whether a given AI may qualify as Civic AI and be admitted near democratic authority.

Remember the Civic Word hypothetical scenario I opened with? Our *Fifty-Gates* would have rejected "Civic Word" because it violates the tests of Illumination, Non-Agency, and Moral-Outsourcing (Gates 7, 4, and 12) - ensuring that project ranking remains a human debate, not a software feature.

No independent decision-making. No civic scoring. No person-level ideological ranking. No hidden authority. And if those boundaries cannot be maintained, the system's proper democratic response is refusal. Because the mission of Civic AI is not power; it is sight.

Not capacity. Not agency. Democratic illumination.

AI near democratic authority must be designed very differently from the systems we're often encouraged to imagine. Many people now hope for a kind of machine savior - a system that will decide faster, optimize better, persuade more effectively.

But optimization is not wisdom. A system optimized for speed can rush toward a cliff of unintended consequences and carry trusting institutions over the edge.

Democracy does not need a machine savior. It needs better light.

To understand the relationship between conscience and tools, let's quickly consider three American luminaries. Each acted from conscience. And in each case, a tool extended human reach:

- Abraham Lincoln's speeches traveled through newspapers and telegraph networks.

- Harriet Beecher Stowe's mass-produced novel, *Uncle Tom's Cabin*, forced readers to see the human reality of slavery.

- Booker T. Washington helped to inspire the Rosenwald Schools Movement, which produced nearly 5000 schoolhouses that expanded educational opportunity for African American children across the rural South.

In each case, human conscience moved first. The tool followed and amplified the human intervention. The tool extended reach, but was never given authority.

That same asymmetry should guide Civic AI today. Humans must pursue civic ends. Systems must be designed so they cannot quietly undermine those ends. Human beings honor democratic principles through intentional action. Civic AI honors democratic principles through designed constraint.

Both may align toward the same civic horizon - but from very different roles. The *Fifty-Gates Democratic Permissions Threshold* instrument makes those boundaries explicit. It treats any Civic AI as a proposed public institution that must earn permission to exist near public power. It insists on non-agency, human moral primacy, reversibility, and anti-weaponization, and it protects pluralism and expressions of conscience. Fail any one of the fifty democratic gates and the AI is disqualified from public service.

The complete *Fifty-Gates* instrument can be accessed at Google Books. I invite all of you to critique, adapt, and extend these ideas in your own work. The *Six Rings* concisely map how digital developments are overwhelming Americans - both individually and across civic life.

All of us - academics, policymakers, AI developers, and citizens - must recognize that agentic AI operating near governance poses a direct threat to 2500 years of democratic evolution.

But a different path is available: non-agentic, democratically-certified Civic AI.

My *Civic Digital Visibility First* principle, together with the *Fifty-Gates Democratic Permissions Threshold*, attempts to reframe the alignment challenge and to systematically test every new AI system before it approaches public power.

Yet civic digital invisibility will continue to remain a structural threat to us all - whether we choose to recognize it or not.

We are at a critical juncture in the emerging field of Civic AI. Let us redefine "Civic AI" clearly - while we still can. Any system that approaches public power without first treating democratic principles as non-negotiable limits is not Civic AI at all; it is simply powerful AI operating near democracy. Civic AI, by definition, begins with constitutional constraint: non-agency, human moral primacy, and refusal-first admissibility. Anything less may be technologically impressive, but it does not qualify for a place near public authority. In other words, *Civic Digital Visibility First* is a governing principle of Civic AI - not a subtype of it.

In the years ahead, human beings - our future Lincolns, and Stowes and Washingtons - will continue to paint the civic canvas. But the role of Civic AI is not to hold the brush.

Edith Wharton once wrote: "There are two ways of spreading light: to be the candle, or the mirror that reflects it."

Let human conscience remain the candle. Then Civic AI may serve as the mirror that reflects that light into the civic world.

Thank You.